

A Framework for Measuring Semantic Drift in Ontologies

Thanos G. Stavropoulos, Stelios Andreadis, Marina Riga, Efstratios Kontopoulos,
Panagiotis Mitziias, Ioannis Kompatsiaris

Centre for Research & Technology Hellas

6th Km Charilaou - Thessaloniki

57001 Thessaloniki, Greece

(+30) 2311 257738

{athstavr, andreadisst, mriga, skontopo, pmitziias, ikom}@iti.gr

ABSTRACT

Semantic drift is an active field of research, aiming to identify and measure changes in ontologies across time and versions, closely related to ontology evolution. However, practical and widely adopted methods that are directly applicable to Semantic Web constructs have yet to emerge. Building upon and extending existing work, this paper presents a framework for measuring semantic drift in ontologies across time or multiple versions, using text and structural similarity methods to provide valuable insights. Its applicability and usefulness are validated through a proof-of-concept scenario in Digital Preservation, where long-term insights about change are crucial, to track drift across a decade's worth of real-world digital media data.

CCS Concepts

- Information systems → Semantic web description languages
- Theory of computation → Semantics and reasoning

Keywords

Semantic drift; concept drift; semantic change; digital preservation; ontologies.

1. INTRODUCTION

Evolving semantics, also referred to as *semantic change*, is an active and growing area of research that observes and measures the phenomenon of change in the meaning of concepts within knowledge representation models, along with their potential replacement by other meanings over time. In the *Semantic Web* (also known as Web 3.0), the representation of the underlying knowledge is typically assumed by ontologies. Thus, it can be easily perceived that semantic change can have drastic consequences on the use of ontologies in Semantic Web and Linked Data applications. In this setting, semantic change relates to various lines of research such as ontology change, evolution, management and versioning [6] [21].

Several ambiguous terms are entailed in the field, bearing slightly different meaning. Namely, *semantic drift* usually refers to ontologies, or other formalisms relying on Semantic Web technologies. *Concept drift*, on the other hand, may refer either simply to the concepts within these constructs or even to machine learning concepts to be learned, referring to an entirely different field. Other terms include semantic shift and decay bearing slightly different meanings [13] [15] [19].

Ontology evolution is another relevant field, which can be defined as the process of an ontology changing in terms of size, content and management in order to accommodate dynamic changes and knowledge interchange in industrial and academic applications. This phenomenon mandates the need for an efficient monitoring and management process [15]. However, well-established, concrete methods to actually measure and compare semantic drift

in ontologies across two or more versions or time in the long-term are yet to emerge.

In an effort to address this issue, this paper initially presents a background study disambiguating terms and clarifying existing approaches regarding change, drift and shift measurement. Then, our framework is presented for concretely measuring semantic drift across multiple versions of ontologies. Our work builds upon existing work in the field of concept drift [18] and extends the proposed methods, implementing them inside an open and reusable software solution. Through this, the framework provides a method to monitor evolving semantics, as a vehicle to measure and manage ontology change. Furthermore, a realistic real-world application in a Digital Preservation scenario is presented, demonstrating our tool's applicability and usefulness in a field where long-term change insights are crucial.

The paper is structured as follows: a background study regarding semantic drift is presented in the next section, so as to disambiguate relevant terms. Section 3 presents related work, i.e. studies that aim to provide means and tools for measuring drift. Section 4 presents the proposed framework and the tools developed for this purpose, followed by Section 5 presenting a realistic, proof-of-concept application of the framework in the field of digital preservation. The paper is concluded with directions for future work and conclusive discussions.

2. BACKGROUND

This section presents an overview of the scientific background on semantic drift, focused on disambiguating the various terms that this field entails. Previous works describe semantic drift using miscellaneous terms. While many times such terms are used interchangeably as synonyms, their use might intentionally imply subtle differences in meaning. In this overview, we treat each term as a different research topic. For each topic we list the number of works that focus on it, either directly or indirectly, along with the meaning they attach to it. Notably, we consider that an existing study directly targets a topic when it provides definitions and methods for measuring metrics of change, and indirectly when they consider it a secondary target, or simply use it as a synonymous term for their main, direct target.

Semantic change refers to the extensive revisions of a single ontology or the differences between two ontologies and can, therefore, be associated with versioning. Semantic change occurs when the internal structure of a concept in two ontologies is different [16]. On the contrary, an isomorphic change refers to the structure being unchanged while the names might have been altered. Others consider semantic change to occur when an ontology revision presents so many alterations that it can be reformed as a new conceptualization, with its own identity [6]. In all cases, the ontology authors are responsible to decide whether

this will occur or the ontology will continue to represent the same conceptualization.

Semantic drift refers to how the features of ontology concepts gradually change, as their knowledge domain evolves. Alternatively, semantic drift refers to the ability of concepts to be re-interpreted by different user communities or in a different context, introducing a risk for them to lose their rhetorical, descriptive and applicative power [20]. It can also be defined as the gradual change of a concept's semantic value, as it is perceived by a community. It can be characterized as intrinsic or extrinsic, depending on whether a concept's semantic value is changed with respect to other concepts in the ontology or to the phenomena it describes in the real world.

Drift can also be classified as either non-collective, inconsistent or consistent collective [4]. If a concept is exposed to extrinsic, but not intrinsic drift, it means that the whole ontology is undergoing a consistent, collective drift that may not necessitate any changes to it. On the other hand, no extrinsic drift together with substantial intrinsic drift means that relationships of concept to others may no longer be correct, even though the concepts themselves have not changed their meaning. In cases of both extrinsic and intrinsic drift we are dealing with inconsistent collective drift rendering the ontology no longer valid.

Concept drift is defined as a change in the meaning of a concept over time, possibly also across locations or cultures, etc. It often refers to a problem in the field of data mining and machine learning, when learned models lose their predictive power over time [18]. In this direction, semantic and concept drift can be considered as two entirely separate fields. Concept drift can refer to the abrupt parameter value changes that occur in data mining, while semantic drift is the language-related version of the same phenomenon [20]. However, previous works have begun bridging the gap between the two fields, by applying notions from concept drift in data mining, such as label, intension and extension, in semantic drift as means for measuring it. More works have adopted this consideration of three types of drift: concept *label drift*, *intensional drift*¹ and *extensional drift* (i.e. a change of meaning that affects the extension of a concept) [9].

Concept shift refers to the subtle changes in meaning of related concepts over time. It can be studied by using chains of extensional, i.e. instance-centric, mapping that represent those subtle changes [19]. Concept shift often occurs in the course of evolution so that the actual meaning of concepts better represent the structure of the real world. While some shifts of concept meaning are performed explicitly, they can also be implicit, through changes in other parts of the ontology, e.g. in properties [15]. The term “topic shift” can also describe the same phenomenon [18].

Concept change refers to the broad variety of adaptations and alterations that can occur for a concept in an ontology. Such changes can be either conceptual (e.g. changing concept relations), specification or representation [21]. Concept change can be tracked by investigating obsolete concepts that have changed name, but maintained their identifiers and a history of changes that can later be examined [17].

¹ In this paper, we will refer to “*intension*” (not to be confused with “*intention*”) as “*the internal content of a term or concept that constitutes its formal definition*” – definition taken from Encyclopedia Britannica online:

<https://www.britannica.com/topic/intension>

Concept or ontology versioning refers to building, managing and providing access to different versions of an ontology [21]. Another definition is that versioning methodology provides users of the ontology variants with a mechanism to disambiguate the interpretation of concepts [6]. It is also linked (but not identical) to ontology evolution by the ontology and database engineering communities, as both research fields aim to represent change and handle different variants of ontologies [21]. However, one of the differences between them is that ontology evolution concerns changing an existing ontology while maintaining consistency, whereas ontology versioning follows a “copy-first” strategy where changes are effected in a new, duplicate version of an ontology [5].

Semantic decay refers to the declination of semantic richness of concepts. The amount of facts that can be inferred from a concept, within the context of Linked Data and a particular dataset, has been proposed as a measure of this richness and thus semantic decay. By using this metric, it has also been proved that the more a concept is reused, the less semantically rich it becomes [13].

Overall, concept drift, concept change and concept shift are closely related, but still connected to semantic drift, while semantic decay remains a self-contained field of study. This overview reveals not only the proximity of various terms, but also their penetration and adoption. For instance, the most popular terms and fields are concept change, concept drift and concept shift, with the least popular being semantic decay, topic shift and topic drift. However, it should be noted that some of the terms (e.g. semantic decay) seem to have been introduced only very recently when compared to others.

According to this background study, we focus on the field of semantic drift and use this term to describe changes in the concepts of ontologies across time or editions. While, the term concept shift is quite close, it most usually refers to concepts outside an ontology, where established methods exist to measure it. We aim to exactly extend methods as those proposed in [18], bridging the gap between semantic and concept drift, theory and practice.

3. RELATED WORK

This section examines existing work that aims to provide means for measuring semantic drift. Measures of semantic richness of Linked Data concepts have been investigated in [13], proving that increasing reuse of concepts, decreases its semantic richness. Other studies have examined change detection between two ontologies at a structural or content level [16]. Concept drift has been measured either by clustering while populating ontologies [1] or by applying linguistic techniques on textual concept descriptions [4]. A vector space model by random indexing has been utilized to track changes of an evolving text collection [20]. A strategy to represent change has been based on ontology evolution [15]. However, most of these techniques are not directly applicable to Semantic Web constructs or present limited statistical data.

One of the most appealing solutions transfers notions applied in machine learning concept drift to semantic drift [18]. The notions of label, extension and intension are not only introduced, but also further defined in ontology terms which renders them highly applicable. Namely, label drift refers to class descriptions in text, intension refers to properties and extension to class instances. Building upon previous work [19], Wang et al. [18] introduce the identity-based and morphing-based approaches, which refer to whether the chain of corresponding identities of classes across ontology versions is known or unknown. This aspect examines

whether the identity of concepts is persistent (and known) across versions or, most often, time. Much philosophical debate is invested in how and by which properties one can identify a concept across time and how this can be formalized [3]. Some approaches utilize the notions of perdurance and endurance, as defined in [2], so as to seek identity. Specifically, looking at rigid, properties that have to be persistent for all instances of a concept can be used to identify entities [10].

Further works have followed up, focusing on the extensional drift aspect of statistical data [9]. Given the high applicability of the method, we adopted it in this framework with the intention to extend and further define methods and tools for its wider adoption in the Semantic Web community.

4. MEASURING SEMANTIC DRIFT

This section presents a framework for measuring semantic drift, i.e. a change in the meaning of a concept over time, location, culture, etc. The underlying methods stem from previous work in the field of concept drift [18]. In these previous studies, highly applicable notions and metrics for measuring concept drift in the context of data mining have successfully been transferred to semantic drift.

In detail, the method to measure concept drift in semantics considers two basic pillars of change: (a) the different aspects of change, and (b) whether concept identity is known or not. The different types of change, reflecting its meaning, include:

- *Label*, which refers to the description of a concept, via its name or title;
- *Intension*, which refers to the characteristics implied by it, via its properties;
- *Extension*, which refers to the set of things it extends to, via its number of instances.

Meanwhile, the correspondence of a concept across versions can be either *known* or *unknown*, resulting in two different approaches for measuring change:

- *Identity-based* approach (i.e. known concept identity): Assessing the extent of shift or stability of a concept’s meaning is performed under the assumption that its identity is known across ontologies. For instance, considering an ontology A, and its evolution, ontology B, each concept of A is known to correspond to a single, known concept of B.
- *Morphing-based* approach (i.e. unknown concept identity): Each concept is pertaining to just a single moment in time (ontology), while its identity is unknown across versions (ontologies), as it constantly evolves/morphs into new, even highly similar, concepts. Therefore, its change has to be measured in comparison to every concept of an evolved ontology.

The contribution of this paper in this area is the adoption, implementation and extension of these methods, in an open, reusable software solution, which is so far missing. The rest of the section describes our proposed method to measure drift, the dataset synthesized for a proof-of-concept scenario and its results. The future work presented in the corresponding section promises to mend many shortcomings in the field of semantic change by providing an open, domain-independent toolbox.

4.1 Method Description

While the implementation of metrics for the different aspects remains generally applicable, the currently proposed method considers the *morphing-based approach*. Despite several methods

have been proposed to seek identity correspondence across versions [10], they still can be domain or model dependent, mandating for ad-hoc expert knowledge in the form of annotations, user input or using explicit identities. For this initial prototype method to remain as generally-applicable as possible, without prior processing and user input, we follow the morphing-based approach, assuming each concept morphs into a new highly similar one in each version. Drift is, hence, measured as the dissimilarity of two maximally similar concepts in two versions [18].

As a result, the method accepts as input only a set of ontology versions, originating from any domain, ordered according to the course of change, e.g. time or locations. As output, for each concept in each version, the method generates three measurements of change (label, intensional and extensional) against concepts in the next ontology in order. For each version, it also generates the average concept change to the next version, for all concepts and for each of the three types, presenting an overview of concept change or stability across versions.

In detail, in order to measure change, the meaning of each concept at a given point t (e.g. in time) is defined as a set of the three different aspects, as follows:

$$C^t = \langle label_t(C), int_t(C), ext_t(C) \rangle$$

where C^t denotes the meaning of concept C at point t , $label_t(C)$ denotes the label aspect of concept C at point t , $int_t(C)$ denotes the intensional aspect of concept C at point t and $ext_t(C)$ denotes the extensional aspect of concept C at point t .

Furthermore, each aspect can be measured as follows:

$$label_t(C) = \{l \mid \forall \langle C, rdfs:label, l \rangle \in T\}$$

$$int_t(C) = \{i \mid i = \langle C, p, x \rangle \vee i = \langle x, p, C \rangle, p = rdfs:domain \vee p = rdfs:range, \forall i \in T\}$$

$$ext_t(C) = \{x \mid \forall \langle x, rdf:type, C \rangle \in T\}$$

where T is the set of all triples in the ontology version t .

In other words:

- The label aspect is given by the `rdfs:label` of a concept.
- The intensional aspect is a set comprised of the union of all RDF triples with C in the subject or object position of OWL Object Properties or OWL Datatype Properties.
- The extension aspect is defined as the set of all instances of `rdf:type C`.

Overall, label is a string, intension is a set of triples and extension is a set of strings. The given choices on how to resolve and formulate each aspect are made both intuitively and based on [18]. Namely, the label is selected as the most representative title

Table 1. Similarity metrics to measure concept differentiation across ontology versions.

Aspect	Similarity Metric to Measure Differentiation
Label	String similarity with Monge-Elkan
Intension	Jaccard similarity between sets of triples
Extension	Jaccard similarity between sets of strings

description of a concept, intension refers to all properties that involve the concept, equally as subject or object, and the extension refers to the number of instances that belong to this concept.

Based on these definitions and using appropriate similarity metrics, one can measure the change/evolution of aspects across versions of the ontology. Table 1 summarizes the metrics adopted in our approach. While Jaccard similarity for comparing sets is a standard solution in this case, altering string similarity metrics is proposed in this work. Namely, the Monge-Elkan algorithm [12] has been found to optimally suit strings in ontologies which are often written in CamelCase or snake_case, without the need for pre-processing (e.g. splits) [14].

In all cases, each concept of an ontology version at point t_1 is compared to all concepts of the version next in order, at point t_2 , for each of the aspects. Due to the morphing-based approach, where identity and correspondence to a specific concept is unknown, all concepts at t_2 should account for measuring drift. Therefore, in all aspects we consider the average similarity between C in t_1 and all concepts of t_2 so as to estimate drift between the two versions.

For *label drift*, the two strings are compared based on the text similarity algorithm Monge-Elkan, which empirically has shown to be more effective for strings met in ontologies. More precisely, in the morphing-based approach, the *label drift* of a concept C between versions t_1 and t_2 is defined as the average of Monge-Elkan text similarity between C in t_1 and all concepts of t_2 :

$$label_{t_1 \rightarrow t_2}(C) = \frac{\sum_{i=1}^{n_2} MongeElkan(label_{t_1}(C), label_{t_2}(C_i))}{n_2}$$

where $label_{t_1 \rightarrow t_2}(C)$ is the label drift of C between versions t_1 and t_2 , $label_{t_1}(C)$ is a string representing the label aspect of C at point t_1 and n_2 is the total number of concepts in t_2 .

In order to measure the similarity of two sets, we deploy the Jaccard similarity, which is defined as follows:

$$Jaccard(A, B) = \frac{A \cap B}{A \cup B}$$

where A, B are two sets of items. Based on that, we define the *intensional drift* of a concept C between versions t_1 and t_2 as the average of the Jaccard similarities between C in t_1 and all concepts of t_2 . This is defined as:

$$int_{t_1 \rightarrow t_2}(C) = \frac{\sum_{i=1}^{n_2} Jaccard(int_{t_1}(C), int_{t_2}(C_i))}{n_2}$$

where $int_{t_1 \rightarrow t_2}(C)$ is the intensional drift of C between versions t_1 and t_2 , $int_{t_1}(C)$ is a set of triples representing the intension of C at point t_1 (properties) and n_2 is the total number of concepts in t_2 .

Similarly, we define the *extensional drift* of concept C between versions t_1 and t_2 as the average of the Jaccard similarities between C in t_1 and all concepts of t_2 .

$$ext_{t_1 \rightarrow t_2}(C) = \frac{\sum_{i=1}^{n_2} Jaccard(ext_{t_1}(C), ext_{t_2}(C_i))}{n_2}$$

where $ext_{t_1 \rightarrow t_2}(C)$ is the extensional drift of C between versions t_1 and t_2 , $ext_{t_1}(C)$ is a set of strings representing the extension of

C at point t_1 (instances) and n_2 is the total number of concepts in t_2 .

Finally, the *whole drift* of concept C between versions t_1 and t_2 , is defined as the average of label, intensional and extensional drift between the same versions:

$$whole_{t_1 \rightarrow t_2}(C) = \frac{label_{t_1 \rightarrow t_2}(C) + int_{t_1 \rightarrow t_2}(C) + ext_{t_1 \rightarrow t_2}(C)}{3}$$

4.2 Implementation

The above method was implemented as a software tool, in order to reproduce the results and apply the methods in multiple occasions, encouraging domain-independent semantic drift research. The current version of the software tool is implemented as a command-line cross-platform application, using Java. The OWL-API library² was used to handle RDF/OWL operations, while the Simmetrics library³ provided the implementation of Monge-Elkan text similarity measure. The implementation part shows much room for improvement and extension in future work, as discussed towards the end of this paper, promising to contribute significantly to the Semantic Web and semantic drift areas. The developed software, along with the sample dataset described in the next section, are available online⁴ under Apache V2 license.

5. AN APPLICATION IN DIGITAL PRESERVATION

This section presents a proof-of-concept application scenario in the field of digital preservation. This realistic scenario serves as a means for validating the applicability of the framework in real-world conditions. For its purpose, a dataset was synthesized using real information for digital media spanning across a decade. Consequently, the methods of the framework were applied, yielding interesting insights regarding semantic drift across time, which were otherwise inaccessible.

5.1 Dataset

In order to validate the proposed approach and apply the methodology, a realistic dataset was synthesized, by extending a domain ontology representing *Software-based Art (SBA)*, developed within the PERICLES FP7 project⁵. More comprehensive descriptions of the project's domain ontologies modelling the Art & Media (A&M) domain can be found in [8] [11] [7].

The synthetic dataset is comprised of SBA ontology versions across time, modelling the evolution and drift of three relevant concepts, namely the concepts of *Computer-based (CB)*, *Mixed-Media (MM)* and *Software-based (SB)* art. The dataset may be synthetic, but is still realistic, as it was based on an internal report by PERICLES partner Tate Galleries, London⁶ that describes the changes to Tate's cataloguing of 8 SBA artworks in the period 2003-2013.

² <http://owlapi.sourceforge.net/>

³ <https://github.com/Simmetrics/simmetrics>

⁴ SemaDrift Library source code online:

<http://mklab.itit.gr/project/semadrift-measure-semantic-drift-ontologies>,

hosted at MKLab tools: <http://mklab.itit.gr/results/tools>

⁵ <http://www.pericles-project.eu/>

⁶ <http://www.tate.org.uk/>

Table 2. Classification of SBA instances to the three concepts in the ontology. The left-most column contains the cataloguing IDs of the artworks.

	2003	2004	2006	2007	2008	2010	2011	2012	2013
T11812	CB	SB	SB						
T12656			CB	MM	MM	MM	MM	SB	SB
T13251				MM	MM	MM	SB	SB	SB
X26607					MM	MM	MM	MM	MM
T13348					SB	SB	SB	SB	SB
T13645						CB	CB	MM	MM
L02995						MM	MM	MM	MM
X43413								MM	MM

Table 2 illustrates the classification of the eight SBA artworks (left-most column) in the three concepts per year (rest of columns). Thus, the dataset contains a total of 9 semantic models (ontologies) for this period, one model per year, excluding the years when no changes in the cataloguing occurred.

5.2 Results

The proof-of-concept use case to measure semantic drift in the A&M domain was performed by feeding the extended SBA ontology versions, ordered by year, to the developed software tool. The output is presented here, starting from morphing chains for each of the concepts, showing their interrelations in-depth and then an overall graph showing the different measures of stability across versions. Concept stability is measured as similarity, in the range of 0 (completely disjoint) to 1 (entirely identical label/properties/instances), according to label/ intensional/ extensional drift respectively.

Initially, morphing chains show in detail concept similarity for each aspect, and how concept meanings migrate from one concept to another, each year, from 2003 to 2013. Inspecting the label aspect shown in Fig. 1, it is apparent that the highest similarity measure holds between concepts with the same name across versions, demonstrating stability.



Fig. 1. Morphing chains for the label aspect.

Likewise, the intensional aspect, shown in Fig. 2, demonstrates equal stability, as properties do not vary significantly across versions.

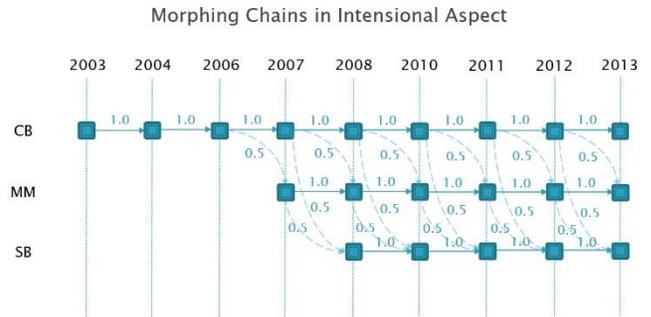


Fig. 2. Morphing chains for the intensional aspect.

On the contrary, the extensional aspect, shown in Fig. 3, demonstrates variations from version to version, with the most significant one being the complete migration of the CB concept partially to MM and to SB concepts, due to its instances shifting type. Actually, this development also coincides with a Tate policy introduced in 2011, according to which the “Computer-based” term was officially abandoned as not sufficiently descriptive and would no longer be used for characterizing the artworks. This demonstrates the capability of the developed methods to interpret to some extent the observed drifts.

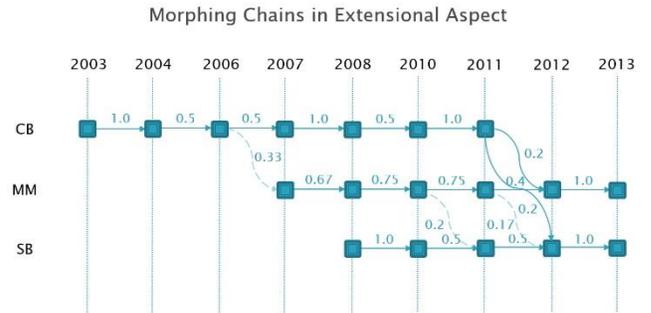


Fig. 3. Morphing chains for the extensional aspect.

Finally, as seen in Fig. 4, the whole aspect depicts these changes in the greater scale, reflecting stabilities (due to label and intension) and some instabilities (due to extension).



Fig. 4. Morphing chains for the whole aspect.

Averaging each aspect for all concepts per version reveals concept stability over time as shown in Fig. 5. This revealing graphic representation clearly shows at a glance that:

- The label aspect is the most stable aspect, followed by intension, since labels and properties remain quite constant in the sample dataset.

- The extensional aspect is the least stable, as all instances of Computer-Based type are eventually evolved into Mixed-Media or Software-Based, as already discussed.
- Stability is reduced in all aspects during 2003-2008, as the ontology is enriched with new concepts.

After looking at concept stability across versions, the Computer-Based concept appears to be the most stable one in all aspects. Notably, we cannot technically average stability for a concept across versions, since in the morphing-based approach their correspondence is unknown. However, we might still measure stability at each given pair, where the Computer-Based concept always ranks first. Still, examining the model more closely reveals that this could actually be attributed to its high similarity to the other two concepts and not to stability per se.

Due to the morphing-based approach comparing to all concepts and this particular concept being highly similar to the other few concepts, similarity across concepts can be falsely perceived as stability here, revealing a limitation of the method. In other words, high similarity across concepts can be interpreted as stability. This limitation is not necessarily misleading, as it could be inherently lifted when enriching the synthetic ontology with more concepts. Meanwhile, this is also a pertinent feature to the morphing-based method, which is dominated by uncertainty, while in an identity-based method the issue would disappear.

6. FUTURE WORK

Future research directions aim to broaden the scope of domain-independent, open tools enabling the Semantic Web community and disseminating the achieved results. So far the core methods developed for calculating drift measures based on ontology evolution have focused on the morphing-based approach, due to its generality and its low requirements from the user (i.e. the ordered set of ontology versions). On the other hand, an identity-based method should also be implemented in the future, as it entails far less uncertainty, giving a much clearer picture of concept stability and drift insights. However, it requires user input to indicate the correspondence of a concept across versions, either through metadata or a graphical user interface (GUI) for user interactions.

The methods themselves can always be enriched with more efficient similarity metrics as done in the current morphing-based methods. As metrics vary, new insights may emerge stemming from limitations – e.g. some metrics for stability may further require normalization. The structure of the ontology could also be investigated as to how it affects drift measures. For example, drift in a single concept could be propagated proportionally to its parent and children nodes using spreading activation.

Furthermore, there are many improvements to implement for both approaches. While the core morphing method is complete, it should be accompanied by a GUI to input basic values such as to indicate file input, order and obtain results graphically, such as those presented in this section. Meanwhile, after implementing the identity approach as well, a GUI will be an even greater facilitator, allowing the user to connect corresponding concepts across versions using graphical means. The tools are planned to be implemented as both standalone cross-platform applications (using JavaFX⁷), or even as Protégé⁸ plugins. The latter being a

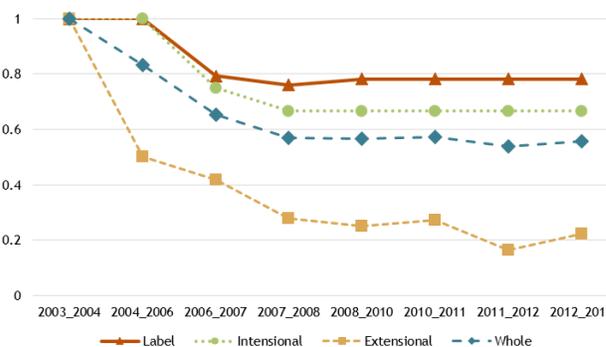


Fig. 5. Average concept stability across time, for each aspect.

very popular and versatile platform in the community will greatly accelerate adoption and dissemination efforts.

7. CONCLUSION

This paper presented a framework for measuring semantic change in terms of semantic concept drift. State-of-the-art methods have been adapted and optimized, measuring label, intensional, extensional and whole (total) drift, inspired by methods in the field of Machine Learning, and following the generic, morphing-based approach. The proposed methodology has been implemented as a domain-independent, cross-platform software tool that will help stimulate research in the area and disseminate the generated results. Consequently, a proof-of-concept experimentation has been performed, by synthesizing a realistic dataset from Art & Media reports from Tate Gallery, showing concept drift in terms of morphing chains, aspect measures and concept stability across time (from 2003 to 2013). The tool shows promise to be extended with more methods and a GUI to facilitate adoption.

Regarding limitations, an issue arises when considering the concept stability measure as presented in Section 5. It seems that the Computer-Based concept is the most stable one, but actually it ranks first because of its high similarity to the other two concepts. This is a feature pertinent to the morphing-based method, dominated by uncertainty and lifted as the ontology grows in size, beyond a synthetic dataset. Notably, it would also be extinct when using an identity approach at the cost of manual labour to annotate the corresponding concepts.

8. ACKNOWLEDGMENTS

This project has received funding from the European Union's Seventh Framework Programme for research, technological development and demonstration under grant agreement no. 601138.

9. REFERENCES

- [1] Fanizzi, N., Amato, C. and Esposito, F. 2008. Conceptual Clustering: Concept Formation, Drift and Novelty Detection. *The Semantic Web: Research and Applications, 5th European Semantic Web Conference, ESWC 2008, Tenerife, Canary Islands, Spain, June 1-5, 2008, Proceedings* (2008), 318–332.
- [2] Gangemi, A., Guarino, N., Masolo, C., Oltramari, A. and Schneider, L. 2002. Sweetening ontologies with DOLCE. *Knowledge Engineering and Knowledge Management: Ontologies and the Semantic Web, Lecture Notes in Computer Science, vol. 2473*. (2002), 223–233.

⁷ JavaFX - <http://www.oracle.com/technetwork/java/javafx>

⁸ Protégé - <http://protege.stanford.edu/>

- [3] Guarino, N. and Welty, C. 2000. A Formal Ontology of Properties. *Knowledge Engineering and Knowledge Management Methods, Models, and Tools*. Springer Berlin Heidelberg. 97–112.
- [4] Gulla, J., Solskinnsbakk, G. and Myrseth, P. 2010. Semantic Drift in Ontologies. *Webist. 2*, (2010), 13–20.
- [5] Klarman, S., Hoekstra, R. and Bron, M. 2009. Versions and applicability of concept definitions in legal ontologies? *CEUR Workshop Proceedings* (2009).
- [6] Klein, M. and Fensel, D. 2001. Ontology versioning on the Semantic Web. *Proc of the Int Semantic Web Working Symposium SWWS. 31*, (2001), 75–91.
- [7] Kontopoulos, E., Riga, M., Mitzias, P., Andreadis, S., Stavropoulos, T., Lagos, N., Vion-Dury, J.-Y., Meditskos, G., Falcão, P., Laurenson, P. and Kompatsiaris, I. 2016. Ontology-based Representation of Context of Use in Digital Preservation. *1st Workshop on Humanities in the Semantic Web - WHiSe, co-located with the 13th Extended Semantic Web Conference (ESWC 2016), At Heraklion, Crete, Greece* (2016), CEUR Workshop Proceedings, Vol-1608, pp. 65–72.
- [8] Lagos, N., Kontopoulos, E., Riga, M., Mitzias, P., Meditskos, G., Waddington, S. and Laurenson, P. 2015. Designing for inconsistency—the dependency-based PERICLES approach. *East European Conference on Advances in Databases and Information Systems* (2015), 458–467.
- [9] Meroño-Peñuela, A., Guéret, C., Hoekstra, R. and Schlobach, S. 2013. Detecting and reporting extensional concept drift in statistical linked data. *CEUR Workshop Proceedings* (2013).
- [10] Meroño-Peñuela, A. and Hoekstra, R. 2014. What Is Linked Historical Data? Springer International Publishing. 282–287.
- [11] Mitzias, P., Riga, M., Waddington, S., Kontopoulos, E., Meditskos, G., Laurenson, P. and Kompatsiaris, I. 2015. An ontology design pattern for digital video. *Proceedings of the 6th Workshop on Ontology and Semantic Web Patterns (WOP 2015)* (2015).
- [12] Monge, A.E., Elkan, C. and others 1996. The Field Matching Problem: Algorithms and Applications. *KDD* (1996), 267–270.
- [13] Pareti, P., Klein, E. and Barker, A. 2016. A Linked Data Scalability Challenge: Concept Reuse Leads to Semantic Decay. *Proceedings of the ACM Web Science Conference* (2016).
- [14] Stavropoulos, T., Andreadis, S., Bassiliades, N., Vrakas, D. and Vlahavas, I. 2015. The Tomaco Hybrid Matching Framework for SAWSDL Semantic Web Services. *IEEE Transactions on Services Computing*. PP, 99 (2015), 1–1.
- [15] Stojanovic, L., Maedche, A., Motik, B. and Stojanovic, N. 2002. User-driven ontology evolution management. *Knowledge Engineering and Knowledge Management: Ontologies and the Semantic Web*. (2002), 133–140.
- [16] Tury, M. and Bielíková, M. 2006. An approach to detection ontology changes. *Workshop proceedings of the sixth international conference on Web engineering - ICWE '06*. (2006), 14.
- [17] Uschold, M. 2001. Creating, integrating and maintaining local and global ontologies. *Proceedings of the First Workshop on Ontology Learning (OL-2000) in conjunction with the 14th European Conference on Artificial Intelligence (ECAI 2000)*. (2001).
- [18] Wang, S., Schlobach, S. and Klein, M. 2011. Concept drift and how to identify it. *Journal of Web Semantics*. 9, 3 (2011), 247–265.
- [19] Wang, S., Schlobach, S., Takens, J. and Van Atteveldt, W. 2009. Mapping-chains for studying concept shift in political ontologies. *CEUR Workshop Proceedings* (2009), 13–24.
- [20] Wittek, P., Darányi, S., Kontopoulos, E., Moysiadis, T. and Kompatsiaris, I. 2015. Monitoring term drift based on semantic consistency in an evolving vector field. *Proceedings of the International Joint Conference on Neural Networks* (2015), 1–8.
- [21] Yildiz, B. 2006. Ontology Evolution and Versioning. *Vienna University of Technology, Karlsplatz*. October (2006), 28.